

REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-01-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate, reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

0625

ING
r
of

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED	
4. TITLE AND SUBTITLE Building and Using Large Common Sense Knowledge Bases				5. FUNDING NUMBERS F49620-97-1-0315	
6. AUTHOR(S) Kenneth D Forbus					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Computer Science Department Northwestern University 1890 Maple Avenue Evanston, IL, 60201, USA				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL, USAF AF Office of Scientific Research 801 N Randolph Street, Room 732 Arlington, VA 22203-1977 Patricia A Bell 703-696-9730				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				20011220 109	
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unlimited distribution, esp for any purpose of the US Government				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) This project explored the use of qualitative representations, analogical reasoning techniques, and sketching to create and use large-scale domain theories. We developed and demonstrated qualitative models of economic systems relevant to intelligence analysis, and showed how qualitative spatial reasoning could provide geospatial reasoning relevant for reasoning about military courses of action and trafficability. We started producing a new analogical processing technology, based on our earlier cognitive simulation work on analogical reasoning. We developed techniques that enabled realistic-sized situations to be compared, for automatically extracting and extending cases dynamically based on task constraints from a general purpose knowledge base, and how these techniques could be smoothly integrated with first-principles reasoning. We explored sketching as a means of communication between people and software. We developed new techniques for hosting knowledge bases on top of relational databases.					
14. SUBJECT TERMS				15. NUMBER OF PAGES	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE		19. SECURITY CLASSIFICATION OF ABSTRACT	
				20. LIMITATION OF ABSTRACT	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

**Building and Using Large Common Sense
Knowledge Bases**

FINAL REPORT

**Kenneth D. Forbus
Computer Science Department
Northwestern University
1890 Maple Avenue
Evanston, IL, 60201**

1. Abstract of Technical Progress

During this project we

- developed a domain theory for trafficability of ground vehicles.
- developed algorithms for constructing Combined Obstacle Overlays and Combined Factor Overlays automatically from GIS data
- developed a domain theory for flow of economic goods that was used in the Crisis Management Challenge Problem by the KSL knowledge server.
- developed a new approach to multimodal interfaces, with an emphasis on deep conceptual and visual understanding of the sketched material, in contrast with the usual focus on recognition in such systems.
- demonstrated that sketch understanding algorithms could be created to interpret simple drawings involving Army operational symbols.
- developed spatial reasoning techniques that can provide human-like geospatial reasoning, given as input military course of action diagrams.
- developed techniques for hosting large knowledge bases in database systems, thus leveraging commercial database products and their features in service of building larger-scale reasoning systems.
- implemented a Domain Theory development Environment (DTE), to test these ideas and to support our experiments.

2. Objectives

We believe that effectively creating and using large-scale domain theories (10,000 - 100,000 axiom equivalents) will require improving the bandwidth of human-computer interaction in several dimensions: the kinds of concepts used, the reasoning methods, and the interaction media. More specifically,

- At the conceptual level, *qualitative representations* provide an ideal medium for expressing the framework knowledge of the physical world that both provides simple inferences directly and also organizes and orchestrates the use of more detailed knowledge.
- At the level of reasoning methods, *analogical reasoning techniques* based on the best available theories in cognitive science can provide more human-like case-based reasoning systems.
- At the level of interaction media, the use of *sketching* can provide a more natural way for domain experts and software to communicate about the spatial aspects of concepts and situations.

These levels strongly interact. By tackling the right aspects of all three at once, as part of a community that is generating large-scale domain theories, we believe we can generate breakthrough capabilities. By incorporating these capabilities into a development environment, we can transfer this technology to other developers and users. Specifically,

- We began developing *a set of analogical processing systems that can serve as a foundation for a new generation of case-based reasoning systems*. In this project we developed a first-cut *analogy ontology* for integrating analogical reasoning with first-principles reasoning, *automatic case construction* methods to formulate

task-specific cases from the contents of general-purpose knowledge bases, and *dynamic case expansion* techniques to enable larger analogies to be created, by growing case descriptions on demand.

- We began developing *sketching as a tool for knowledge acquisition and refinement*. Our long-range goal is to make sketching as natural a modality for explaining things to machines as it is for explaining things to people. Our approach, which focuses on rich spatial reasoning and conceptual understanding rather than focusing on recognition, explores a very different set of tradeoffs and capabilities than previous work on multimodal interfaces.
- We created a *Domain Theory Development Environment* that uses relational database technology to store general-purpose knowledge bases, and uses a *federated architecture* to integrate specialized reasoning facilities, such as geographic information systems, spatial reasoning software, and our analogical processing systems.

3. Status of Research Effort

Our progress on this project suggests that the techniques of qualitative representation and reasoning that we originally developed for engineering tasks and domains are directly applicable to military tasks and domains. For example,

- We developed a domain theory for trafficability, based on Army Doctrine. This domain theory uses representations of mathematics (including qualitative mathematics) and default reasoning to exploit available information.
- We developed algorithms for constructing Combined Obstacle Overlays and Combined Factor Overlays automatically from Geographic Information System (GIS) data. The fundamental approach is identical to what we did earlier in developing theories of reasoning about mechanical systems (e.g., clocks). That is, spatial reasoning involves two representations. The metric diagram is the equivalent of our perceptual apparatus. Place vocabularies are task-specific qualitative descriptions of shape and space, computed automatically from the metric diagram. In the case of terrain reasoning, the metric diagram is a GIS, and the COO and CFO are two place vocabularies. The place vocabularies are constructed by finding regions where the trafficability parameters are equivalent.
- We developed a domain theory for flow of economic goods that was used in the Crisis Management Challenge Problem. This domain theory used techniques from Qualitative Process theory and compositional modeling to describe flows of goods in the world oil and natural gas economies. Since one of the goals was to demonstrate that the creating of libraries of domain theories is possible, it is notable that our flow domain theory was used by someone else, as part of their integrated system, to answer a number of the questions.

We also made progress on sketch understanding, demonstrating that sketch understanding algorithms that relied on visual reasoning and analogical encoding could be used to understand simple drawings involving Army operational symbols, including retrieving drawings involving similar plans.

We also implemented a Domain Theory development Environment (DTE), to support our experiments. This lightweight environment uses a variation of Hendler's techniques for encoding a knowledge base into a relational data base, providing efficient access to arbitrary propositional representations instead of just frame systems. We also developed a *federated architecture* for reasoning, which allows specialized reasoners to be tightly integrated. The basic ASK mechanism for queries uses dispatches based on the structure of the query to a *reasoning source*, an interface between a special-purpose system and the rest of the reasoner. Any answers provided by the reasoning source are woven into the system's working memory via a logic-based truth maintenance system, so that explanations can include information as disparate as geographic information system data and analogical reasoning.

This project led to significant strides forward in turning our cognitive simulations of analogical processing into tools that can be used on large-scale problems, what we think of as a step on the road to a human-like analogical processing technology. Here is a summary of the impact of these changes:

Before	After
Typical cases 70 propositions, largest used 400 propositions	Can map cases that include thousands of propositions
No uniform mechanism linking analogy to other reasoning systems	Analogy ontology links 1 st principles reasoning and analogical reasoning
Narrow, project-specific knowledge bases	Working with large KB's built by other groups (e.g., Cyc, SAIC's KB)
Cases pre-packaged by hand	Cases can be automatically constructed from general-purpose KB, based on task constraints
Case libraries distinct from general-purpose knowledge	Case libraries consist of dynamically changing subsets of a KB

The details of these improvements are described in the publications below.

4. Publications

The following publications were based wholly or in part on results from this project:

1. Forbus, K., Mahoney, J.V., and Dill, K. 2001. How qualitative spatial reasoning can improve strategy game AIs. AAAI Spring Symposium on AI and Computer Games.
2. Forbus, K., Ferguson, R. and Usher, J. 2001. Towards a computational model of sketching. *IUI'01*, January 14-17, 2001, Santa Fe, New Mexico.
3. Forbus, K. 2000. Exploring analogy in the large. In Gentner, D., Holyoak, K. and Kokinov, B. (Eds) *Analogy: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.
4. Mostek, T., Forbus, K. and Meverden, C. 2000. Dynamic case creation and expansion for analogical reasoning. *Proceedings of AAAI-2000*. Austin, Texas.

5. Ferguson, R. and Forbus, K. 2000. GeoRep: A flexible tool for spatial representation of line drawings. *Proceedings of AAAI-2000*. Austin, Texas.
6. Forbus, K., Ferguson, R., and Usher, J. 2000. Towards a computational model of sketching. *Proceedings of QR-2000*. Morelia, Mexico.
7. Ferguson, R., Rasch, R., Turmel, B., and Forbus, K. 2000. Qualitative spatial interpretation of course-of-action diagrams. *Proceedings of QR-2000*. Morelia, Mexico.
8. Ferguson, R. W. and Forbus, K.D. GeoRep: A flexible tool for spatial representation of line drawings. 13th International workshop on qualitative reasoning (QR99), Loch Awe, Scotland. June 6-9, 1999.
9. Donlon, J.J. and Forbus, K. D. Using a geographic information system for qualitative spatial reasoning about trafficability. 13th International workshop on qualitative reasoning (QR99), Loch Awe, Scotland. June 6-9, 1999.

5. Personnel

Army Interns: CPT James Donlon, CPT Rob Rasch.
 Programmers: George Lee, Donna Furtzie, Jeff Usher
 Graduate Students: Tom Mostek, Jeff Usher, Mike Brokowski
 Undergraduate Students: Cara Meverden,
 Research Scientist: Ron Ferguson

6. Significant events involving project personnel

- Jeff Usher, who started out as a graduate student, decided to become a programmer instead.
- Cara Meverden now works for the US Government.

7. Interactions and Transitions

We had "in-house" informants regarding Army practices concerning battlespace reasoning, in the form of US Army interns. Interns take courses and work on research projects, receiving a Master's degree in Computer Science after two years. We have found that teaching domain experts the necessary skills of knowledge representation and qualitative reasoning to be an effective strategy for creating relevant domain theories.

Our work on automated trafficability reasoning was delivered to a classified US Military customer, after a government contractor did some reengineering of it with our assistance.

Our work on COA sketching has been transitioned to DARPA's Command Post of the Future program. An experiment with our COA software, using Cadet as the back-end to generate a synchronization matrix, was evaluated in an experiment at the Battle Command Battle Lab in Ft. Leavenworth as part of the ICCES program in November 2000. While the number of subjects was too small to draw conclusions about statistical significance, results with military personnel indicated that the ability to sketch doctrinal

COAs quickly and have synchronization matrices produced from them automatically was an exciting new capability.

8. Presentations

The following invited presentations included descriptions of some of the results of this research. (Standard conference paper presentations not included)

- Naval Research Laboratory, Washington, D.C. February, 2001
- University of Buffalo Cognitive Science Program, February, 2001
- NSF Modeling and Visualization Workshop, Washington, D.C. October, 2001.
- MONET Model-based reasoning and qualitative reasoning Summer School, Bertinoro, Italy, May, 2000
- Distinguished Lecturer in Cognitive Science, Stanford University, Palo Alto, California, April 2000
- Chukyo University, Nagoya, Japan, March 2000
- Keio University, Tokyo, Japan, March, 2000
- NSF Spatial Intelligence Conference, Chicago, Illinois, October, 1999
- National Conference on Artificial Intelligence (AAAI-99) Orlando, Florida, July, 1999
- Keynote speaker, Visual and Spatial Reasoning in Design Conference, MIT, June, 1999
- Invited talk, 13th International Workshop on Qualitative Reasoning (QR99), Loch Awe, Scotland, June, 1999
- Microsoft Research, Redmond, Washington, September, 1998
- Keynote speaker, 1st International Workshop on Analogy, Sofia, Bulgaria, August, 1998
- Invited talk, AAAI Spring Symposium on Commonsense causation, March, 1998

9. Consulting and advisory roles

Forbus has been a consultant at Xerox PARC since 1984. The next-generation reasoning engine is being built in collaboration with PARC, with PARC supplying a higher-performance database optimized for knowledge storage, and with Northwestern providing analogical processing technology.

10. Additional Perspective

Our analogy research prior to this project had focused on cognitive simulation: Using programs to model psychological results in analogy and similarity. Modeling psychological phenomena well is important for two reasons:

1. Human abilities are the inspiration for approaches such as case-based reasoning. Unfortunately, engineering approaches to CBR that ignore the psychology of it have so far failed to capture the power and flexibility of human analogical reasoning.

2. We think that using within-domain analogies (what are called *literal similarity* matches) over concrete descriptions in combination with first-principles qualitative reasoning is the key to achieving human-like common sense reasoning. First-principles reasoning expresses what is logically possible, but not what actually occurs or is likely to occur. On the other hand, remembered situations and conservative abstractions based on bodies of experiences provide a guide to what is likely. Similar remembered situations can provide a rapid source of plausible predictions. Experience can also act as a reality check for first-principles reasoning: Logically possible predictions that have never been observed in a familiar domain should be viewed with caution. Thus we suspect the flexibility and lack of brittleness in human common sense reasoning derives from a combination of first-principles and analogical reasoning.

In this project we began turning our cognitive simulations of analogy (SME does matching, MAC/FAC does retrieval) into an analogical processing technology that we think could revolutionize case-based reasoning. Unlike most feature-based CBR systems, SME and MAC/FAC handle structured representations, meaning they can (and have) been used with representations of arguments, proofs, plans, equations, and other complex structured representations. MAC/FAC does not require hand-indexing of representations, which means cases can more easily be added by domain experts. Creating a knowledge-rich, integrated environment for experimenting with this technology is one of our key goals.

During this project we reengineered and extended our existing models of analogical matching (SME) and similarity-based retrieval (MAC/FAC) significantly. This included enabling them to operate over larger examples, by a combination of low-level engineering and dynamic case expansion, and with a wider variety of representational formats, by modifying the structural evaluation algorithm to take reified events into account. One key advance was to create an *analogy ontology* that reifies the ideas of comparison, such as correspondences, mappings, matches, and candidate inferences, which enables the constructs of analogical processing to be treated as first-class entities within reasoning systems. This provides the framework for integrating analogical processing with first-principles reasoning. As one of the first steps in using this framework to explore larger-scale issues in analogical reasoning, we developed novel *automatic case construction methods*, where different types of queries could be used to extract knowledge about a situation, event or entity from a general knowledge base to create a task-specific case. We also started exploring the use of first-principles reasoning to validate candidate inferences. Examples include testing the critical antecedents for critiquing military courses of action via analogy and filtering candidate inferences in making predictions about hypothetical international crises based on historical precedents.

The ability to use our analogical processing systems with large knowledge bases and interoperate with other types of reasoning systems constitutes a fundamental shift in the kinds of questions we can tackle and problems we can address.